Learning and Planning for Mars Rover Science

Tara Estlin, Rebecca Castano, Robert Anderson, Daniel Gaines,
Forest Fisher, and Michele Judd
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena CA 91109
Tara.Estlin@jpl.nasa.gov

Abstract

With each new rover mission to Mars, rovers are traveling significantly longer distances. In some cases, distances are increasing by orders of magnitude from previous missions. This increase enables not only the collection of more science data, but causes a large rise in the number of new and different science collection opportunities. In this paper, we describe the OASIS system, which provides autonomous capabilities for dynamically pursuing these science-collection opportunities during long-range rover traverses. OASIS utilizes techniques from both machine learning and planning and scheduling to address this goal. Machine learning techniques are applied to analyze data as it is collected and quickly determine new science tasks and priorities on these tasks. Planning and scheduling techniques are used to alter the rover's behavior so new science measurements can be performed while still obeying resource and other mission constraints. In addition to describing our system, we also discuss how we are testing OASIS, including the use of Mars rover prototypes and validation using data gathered from expert planetary geologists.

1 Introduction

As planetary exploration continues to increase, the use of robotic vehicles to explore and analyze planet surfaces will also expand. The Mars Pathfinder mission not only demonstrated the feasibility of sending rovers to other planets, but displayed the significance of such missions to the scientific community. The Mars Exploration Rovers (MER) mission is set to launch this year, and will send two new rovers to the Martian surface. Furthermore, additional rover missions are already planned to the red planet, which will provide major leaps in smart, surface laboratory measurements. With each new mission, rovers are able to travel significantly longer distances and collect increasing amounts of valuable science data. However, they must perform this task in unknown environments where unexpected conditions can easily be encountered. The Pathfinder rover traveled approximately 100 meters during its 90 day lifetime [Mishkin, *et al*., 1988]. In contrast, the MER rovers will travel up to 100 meters per day, and future missions will likely continue to extend this measure. Though longer-range traverses enable rovers to explore new territory and collect large volumes of data, they also place increasing demands on operating these missions. Collected images and other science data must be analyzed (typically on earth), and this process must be performed quickly if that data is used to direct additional science measurements by the rover. Furthermore, rover operations for both Pathfinder and MER are handled by manually creating sequences of rover commands on the ground and then uploading them to the rover. This process is very time- and labor-intensive and does not allow for the dynamic adjustment of rover behavior if anything unexpected occurs, including faults and new science opportunities.

This paper describes the Onboard Autonomous Science Investigation System (OASIS) [Castano, *et al*., 2003], which is directed at providing autonomous capabilities for rover science operations during long-range traverses. In upcoming missions, rovers will traverse many kilometers between pre-designated science sites. OASIS was developed to support science data analysis and new science collection during these long traverses. OASIS consists of several modules, including: 1) a data analysis system that uses machine learning techniques to analyze collected data and produce new science collection goals and 2) a planning and scheduling system that dynamically incorporates new science goals into the current rover command sequence and interacts with the onboard control software to achieve this goal. This system is currently being tested on data collected during test operations for the MER mission as well as with the Rocky 8 rover, a research rover built and supported at JPL.

Science data analysis in OASIS is performed using several different machine-learning techniques, which can prioritize acquired science data for downlink to earth and create new science goals for the rover to achieve. This paper concentrates on the latter capability of creating new science goals. More information on prioritizing data for downlink can be found in [Castano, *et al*., 2003]. Three different prioritization methods have been devel-

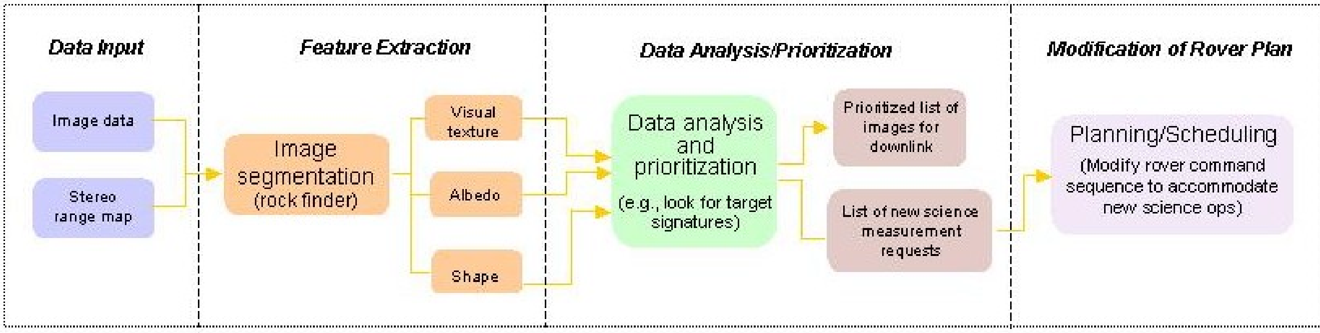

Figure 1. Overview of OASIS system architecture. OASIS consists of three major components: Feature Extraction, Data Analysis/ Prioritization, and Planning and Scheduling.

oped for OASIS. All use extracted rock features to rank rocks in terms of scientific importance. The first technique, *target signature detection*, recognizes pre-specified signatures that have been identified by ground scientists as data of high interest. The second technique, *novelty detection*, identifies unusual signatures that do not conform to the statistical norm for the region. The last technique, *representative sampling*, prioritizes science measurements by ensuring data is collected on representative rocks of the traversed region. These three prioritization methods are used to trigger opportunistic science observations by identifying valuable new science opportunities that, if possible, should be taken advantage of during the rover's traverse.

When science opportunities arise on a traverse, a planning and scheduling system is used to determine the necessary rover activities to achieve the new science goals. Based on an input set of prioritized goals and the rover's current command sequence, the planner generates a modified sequence of activities that satisfies as many new goals as possible while still preserving high-priority activities already in the sequence and obeying resource and other operation constraints (e.g., such as ensuring there will be enough power to complete the day's activities). Our planner uses a *continuous planning* approach, where plans are dynamically modified in response to changing events and goal information. In this approach, the planner continually monitors the execution of commands on the rover and information on resource utilization and current states. It also accepts new science goals as they become available. As information is acquired regarding these items, the planner updates its version of the plan. From these updates, new problems and/or opportunities may arise, requiring the planner to re-plan in order to accommodate the unexpected events.

The remainder of this paper is organized as follows. We begin by presenting the OASIS system, including characterizing the full architecture and presenting a more detailed explanation of the system components. We will then describe our testing plan for OASIS, which includes using Martian data from upcoming missions, as well as robotic vehicles developed at JPL. Finally, we will discuss related work and present our conclusions.

2 OASIS System

The OASIS system architecture is shown in Figure 1. As highlighted in the figure, OASIS is comprised of three major components:

- **Feature Extraction**: Enables extraction of features of interest from collected images of the surrounding terrain. This component both locates rocks in these images and extracts rock properties, such as shape and texture.
- **Data Analysis/Prioritization**: Uses extracted features to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment. This component consists of three different prioritization algorithms, that analyze collected data, prioritize identified rocks, and generate a new set of observation goals to gather further data on rocks which were ranked high priority.
- **Planning and Scheduling**: Enables dynamic modification of the current rover command sequence (or plan) to accommodate new science requests from the data analysis unit. This component uses a *continuous planning* approach to iteratively adjust the plan as new goals and/or faults occur.

OASIS operates in an autonomous fashion where the data analysis system can be seen as driving new science exploration. First, new science data is received by the Feature Extraction component. Currently, we have focused the system on analyzing rocks within \image data, but plan to expand to other types of data, such as spectrometer measurements. Images are broken down by first locating individual rocks in each received image, and second, by extracting a set of rock properties (or features) from each identified rock. Extracted rock properties are then passed to the Data Analysis component of the system.

This component consists of three different prioritization algorithms, which analyze the data by searching for items such as pre-known signatures of interest, which have been identified by scientists on earth, or novel rocks (i.e., outliers) that have not been seen in past traverses.

As shown in Figure 1, this analysis produces two main products. One is a set of prioritized images for transmission to Earth. Currently, spacecraft, such as rovers, can collect significantly more data than can be transmitted to Earth due to communication limits. OASIS ranks images by scientific importance so more valuable images get transmitted first for further analysis on the ground. This paper is focused on the second product, which is a list of new science measurement requests. OASIS uses the output of its three prioritization algorithms to dynamically produce a list of new science measurements that will take advantage of new and interesting data collection opportunities. In current rover missions, images and other science measurements are only sent to earth once or twice during the day. Furthermore, many images cannot be sent at all due to the communication restrictions mentioned above. This setup means that many valuable science opportunities may be lost. One problem is that by the time images are sent from Mars to Earth, analyzed on the ground by scientists, and a new set of measurement requests determined and sent back to Mars, the rover will likely have passed the object of interest. Another problem is the opportunity may never be recognized if the identifying data is never sent to Earth for analysis. By analyzing data onboard, OASIS enables these new science opportunities to be dynamically realized.

New science measurement requests (or goals) are passed to the planning and scheduling module, which produces a modified set of actions in order to achieve as many new science goals as possible, without violating resource or other mission constraints. In current mission operations, rover behavior is directed by manually hard-coding sequences of commands on Earth and then uploading these sequences to the rover. Sequence changes are rarely performed onboard and if something unexpected happens, the rover must contact earth for further instructions. The planning and scheduling component addresses this problem by using a model of rover operations and constraints to dynamically modify the current rover plan in order to accommodate new science goals. This component can also monitor plan execution and continue to modify the rover command plan if other unexpected events or faults occur.

Next, we discuss each of the OASIS components in more detail.

2.1 Feature Extraction

The first step in the OASIS system is analyzing rock features from images taken onboard the rover. As the rover traverses, it takes a series of images to support not only science, but also navigation operations. Images may be taken from several different cameras to capture information on the surrounding terrain for science analysis and/or assist in path planning, obstacle avoidance, etc. Our initial emphasis in OASIS has focused on image analysis and the characterization of surface rocks. Rocks are among the primary features populating the Martian landscape and the understanding of rocks on the surface is a first step leading towards more complex regional geological assessments by a robotic vehicle.

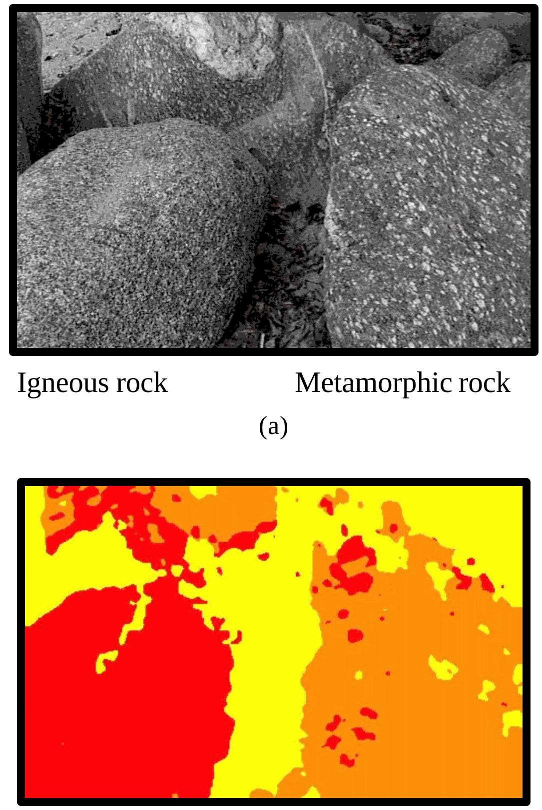

(b)

Figure 2. Examples of visual texture providing information about the geologic texture of rocks. (a) original image (b) image segmented based on texture.

Rocks are located in the images by determining the ground plane from stereo range data, and then producing a height image and level contours for that image. These contours can be connected from peaks to the ground plane to identify rocks in each image [Gor, *et al.*, 2001].

Next, a set of properties is extracted from each rock. Our feature extraction priorities are based upon our knowledge of how a geologist in the field would extract information. Important features to look for and categorize include albedo (an indicator of rock surface reflectance properties), visual texture (which provides valuable clues to mineral composition and geological history), shape, size, color and arrangement of rocks. Currently our system identifies the first three of this set; future work will expand this to cover additional features. Each property or feature is measured using a different technique [Gilmore, *et al*., 2000; Castano *et al.*, 2002]. For instance, visual texture is measured by computing gray-scale intensity variations at different orientations and spatial frequencies within the image. Figure 2 shows visual texture information produced from one sample image.

2.2 Data Analysis and Prioritization

The second step in the OASIS system is to use the extracted features to prioritize rocks. Three prioritization techniques are used, which employ different machine learning methods. The results from this analysis are then used to identify rocks that should be further analyzed and produce a new set of science measurement goals.

Key Target Signature

The first prioritization technique, key target signature, enables scientists to efficiently and easily stipulate the value and importance of certain features. Scientists often have an idea of what they expect to find during a rover mission and/or are looking for specific clues that reflect signs of life or water (past or present). Using this technique, target feature vectors can be pre-specified and an importance value assigned to each of the features. Rocks are then prioritized as a function of the weighted Euclidean distance of their extracted features from the target feature vector.

Novelty Detection

The second prioritization technique, novelty detection, detects and prioritizes unusual rocks that are dissimilar to previous rocks encountered. We have looked at three different learning techniques for novelty detection. First, we have developed a distance-based k-means clustering approach, in which a set of rocks are clustered and any new rock that is a great distance from any of the cluster centers is considered novel. In the second method, the probability density over the feature space for a set of rocks is approximated using a Gaussian mixture model. The novelty of a new rock is inversely proportional to the probability of that rock being generated from the learned mixture model. The third method uses a discrimination based kernel one-class classifier approach. We treat all previous rock data as the "positive class" and learn the discriminant boundary around that class. Future rocks with features falling outside the boundary are considered novel. An example of detecting a novel rock using data collected from rover field tests is shown in Figure 3.

These three approaches represent the three dominant flavors of machine learning techniques for novelty detection: distance-based, probability-based (i.e., generative), and discriminative. Considering all three types in one hybrid approach allows us to tradeoff their respective advantages and disadvantages.

Representative Sampling

In order to understand the region being traversed, it is important to have information on representative rocks, vs. very interesting or unusual rocks. A region is likely to be populated by several types of rocks, with each type having a different abundance. Thus, a uniform sampling

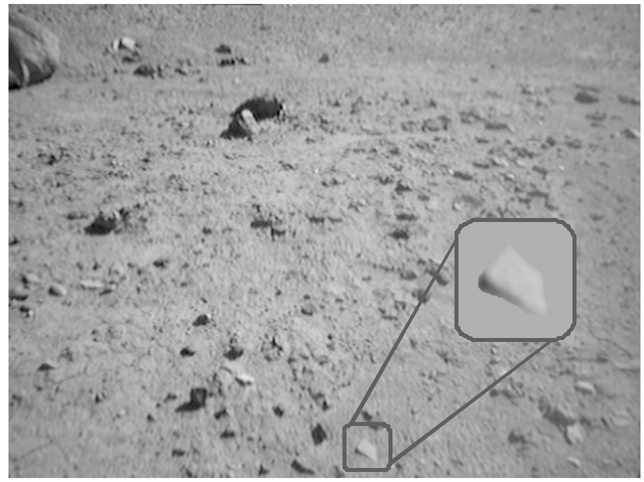

Figure 3: Detection of significant novel rock. The marked rock is a piece of petrified wood that was discovered during rover field tests for the MER mission. This piece of wood was identified as novel by the OASIS system, however was not identified by the remotely located geologists during the rover tests.

will be biased towards the dominant class of rock present and may result in smaller rock classes not being represented at all in the downlinked data.

The third prioritization technique, representative sampling, provides an understanding of the typical characteristics of a region. Rocks are clustered into groups with similar properties and the data is then prioritized to ensure that representative rocks from each class are sampled. To determine the classes, the rock property values are connected together in a series to form a feature vector. A weight is assigned to the importance of each property. Unsupervised clustering is then used to separate the feature vectors into similar classes. We currently use a k-means clustering technique due to its relatively low computational requirements. However, other unsupervised methods could also be employed. For each class of rocks, this technique can find the most representative rock in the class (i.e., the single rock in any image that is closest to the mean of the set) or rank rocks according to this metric.

Science Alert

Using the above determined priorities, the data analysis software can then flag rocks that should be further analyzed and produce a new set of measurement goals to further characterize the identified rocks. We call this capability *science alert*, since it alerts other onboard software that new and high priority science opportunities have been detected. The number of new goals produced by the data analysis software will vary depending on the constraints of the mission. Some missions may want only limited science alert capabilities, and thus new opportunities would only be flagged it they were deemed critical. Other missions may allow onboard analysis to direct a larger portion of planned science measurements.

Science alert may also involve several different levels of reaction. OASIS has been designed so a spectrum of reactions can be supported. The most basic reaction is to adjust the rover plan so that the flagged data is immediately sent back to Earth for further analysis and the rover holds at the current position, delaying other tasks. The next step would likely be to collect additional data at the current site before transmitting data to earth. Further steps include having the rover alter its path to get closer to objects of interest before taking additional measurements and/or scheduling a close contact measurement (such as with a microscopic imager). These operations would provide new data that could not be obtained through image analysis alone. The level of allowed reaction will likely be determined by the constraints and goals of the rover mission. Reaction capabilities may also be allowed to vary over the course of the mission.

2.3 Planning and Scheduling

Once the data analysis software has identified a set of new science targets, these targets are passed to onboard planning and scheduling software that can dynamically modify the current rover plan in order to collect the new science data. This component takes as input the new set of science requests, the current rover command sequence (or plan), and a model of rover operations and constraints. It then evaluates what new science tasks could be added to the current plan while ensuring other critical activities are preserved and no operation or resource constraints are violated.

CASPER Planner

Planning and scheduling capabilities are provided in OASIS by the Continuous Activity Scheduling, Planning and Re-Planning (CASPER) system [Estlin, *et al*., 2002; Chien, *et al*., 2000]. CASPER provides a generic planning and scheduling application framework that can be tailored to specific domains. Its components include:

- An expressive modeling language to allow the user to naturally define the application domain.
- A constraint management system for representing and maintaining domain operability and resource constraints.
- A set of search strategies and repair heuristics.
- A temporal reasoning system for expressing and maintaining temporal constraints.
- A graphical interface for visualizing plans.
- A real-time system that monitors plan execution and modifies the current plan based on activity, goal, state and resource updates.

CASPER employs a *continuous planning* technique where the planner continually evaluates the current plan and modifies it when necessary based on new state and resource information. Rather then consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current rover state, and state projections into the future for that plan. At any time an incremental update to the goals or current state may update the current plan. This update may be an unexpected event (such as a new science opportunity) or a current reading for a particular resource level (such as power). The planner is then responsible for maintaining a plan consistent with the most current information. And since things rarely go as expected, especially during planetary surface operations, the planner stands ready to continually modify the plan.

A plan consists of a set of grounded (i.e., time-tagged) activities that represent different rover actions and behaviors. Activities can be at different levels of abstraction, where low-level activities typically correspond to direct rover commands. For example, a plan typically contains several *traverse* activities that move the rover between different locations in order to visit science targets. Rover state in CASPER is modeled by a set of plan timelines, which contain information on both states, such as rover position, and resources, such as power. Timelines are calculated by reasoning about activity effects and represent the past, current and expected state of the rover over time. As time progresses, the actual state of the rover drifts from the state expected by the timelines, reflecting changes in the world. State updates are relayed back from sensors and the rover control software. As these updates are received, CASPER updates the relevant timeline models with actual state values, resource values, activity completion times, etc. Each of these updates may introduce problems into the current plan, such as a power over-subscription due to a long traverse or an instrument being in the incorrect position to perform a particular science reading. These problems (or plan conflicts) cause CASPER to perform plan modifications to bring the plan back into sync with the current state and set of goals. An example of a plan in the CASPER GUI is shown in Figure 4.

To produce and/or modify a current rover plan, CASPER uses an iterative repair algorithm [Zweben *et al*., 1994], which classifies plan conflicts and attacks them individually. Conflicts occur when a plan constraint has been violated where this constraint could be temporal or involve a resource, state or activity parameter. Conflicts are resolves by using one or more plan modifications such as moving, adding, or deleting activities. One example of a conflict is when a new science activity oversubscribes a resource such as power or memory. Possible resolutions to this conflict might be moving the science activity to a part of the plan that doesn't oversubscribe that resource, deleting the science activity, or moving and/or deleting other contributing activities.

Path Planning

To provide spatial reasoning capabilities to CASPER, we are using a global path-planning module, which provides rover route information to the planner based on a map of the

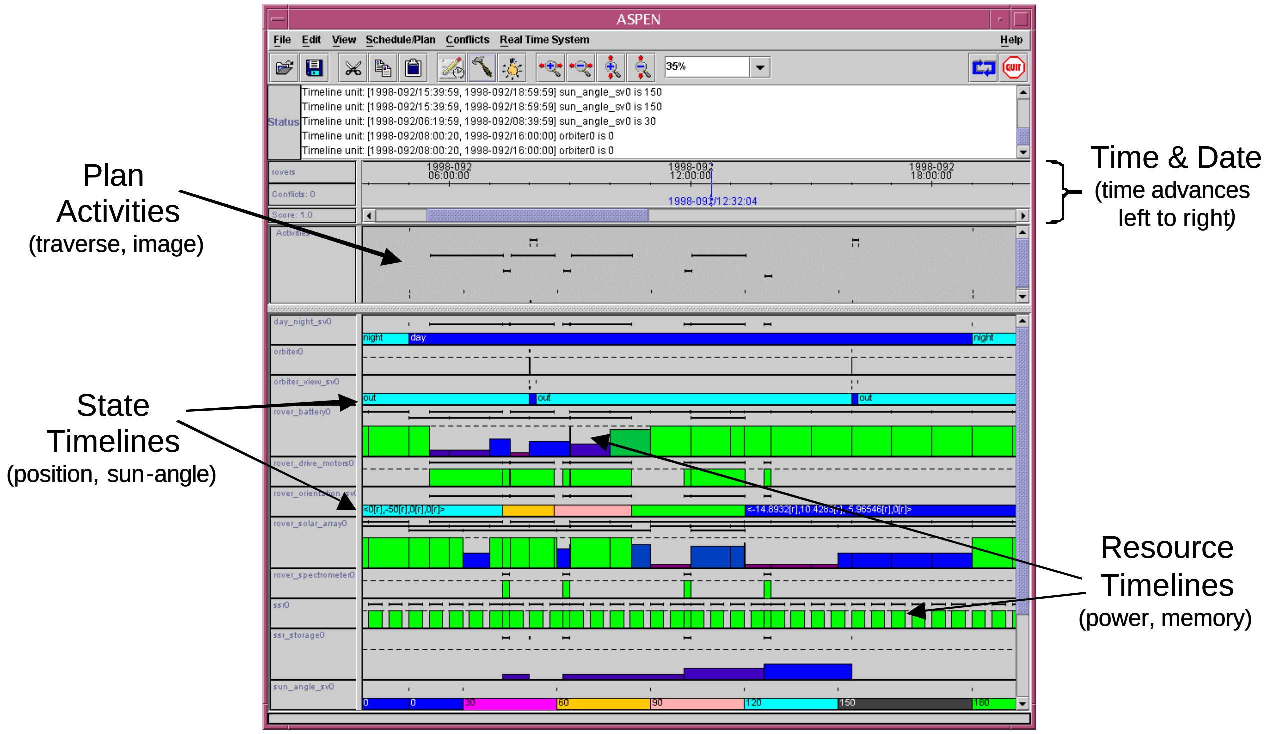

Figure 4: Sample rover plan displayed in planner GUI. Plan activities are shown as bars in upper portion of window, where bars represent the start and end time of each activity. State and resource timelines are shown in bottom portion of screen and show the effects of the plan as time progresses. Time is depicted as advancing from left to right.

rover's surrounding environment. This module is intended to give a global perspective of the rover's anticipated path as opposed to the local perspective that would be considered by obstacle avoidance software. We are assuming that for most rover operations some global map information would be available through orbital or descent imagery, or from panoramic imagery generated onboard the rover itself. We are also assuming this map information map be incomplete and certain terrain features and/or obstacles may be missing.

CASPER queries the path planner for two main pieces of information. The first piece is estimated distances between science target or other designated traverse waypoints. The second piece is a list of intermediate-waypoint coordinates that can be used to direct the rover's traverse to a particular targets. Path-distance information is used by the planner to estimate duration and power required for rover traverses between targets. Intermediate waypoints are used to track the rover's progress during a traverse. To provide path planning information to our system, we are currently using the D* path planner, which produces paths in partially known or changing environments using an efficient and optimal algorithm [Stentz, 1994].

3 System Testing

We are in the process of testing the OASIS system using data gathered during rover field tests for upcoming missions as well as using several JPL research rovers in the JPL Mars Yard.

The data analysis component is currently being tested using a suite of image data collected during rover field experiments performed in Flagstaff, AZ. (These field experiments were done in preparation for the upcoming 2003 MER rover mission.) One of the primary goals of using this data to test OASIS is to not only test our system on realistic data, but to also ensure that the prioritizations our algorithms produce are comparable to those made by planetary geologists. Our approach for testing is to gather sample prioritizations from expert planetary geologists on various collections of images. Expert rankings are input using a web-based application that enable experts from across the country to easily prioritize images and add explanations for their decisions. We are using statistical methods to combine these expert prioritizations and compare them with the prioritizations produced by our algorithms.

The planning component has already been used in several tests [Estlin, *et al*., 2002] using two JPL rovers, Rocky 7 and Rocky 8, which are pictured in Figure 5.

Figure 5: JPL Rocky 7 and Rocky 8 rovers

Rocky 7 is approximately the same size and mass as the 1997 Mars Pathfinder rover, Sojourner. It employs a rocker-bogie six-wheel configuration, and is a partially-steered vehicle, where it only has steering capability on two corners. In contrast, Rocky 8 is roughly an order of magnitude larger than Rocky 7 and is similar in size to the twin MER rovers, set to launch later this year. Rocky 8 also employs a rocker-bogie six-wheel configuration, however it is a fully-steered vehicle with all-wheel drive and all-wheel steering.

The planner was used to produce an initial rover plan based on a set of science objectives (e.g., perform an image at location A, perform a spectrometer read at location B, etc.) and to dynamically modify that plan when unexpected events occurred during execution (e.g., more power was required for a traverse or science activity than originally estimated). Tests were performed in the JPL Mars Yard. The initial plan contained 53 different activity instances and took the planner 3.7 seconds to construct. The planning horizon for these tests was 4 hours. Replans took an average of 7 seconds.[1] During these tests the planner interacted with the rover control software in several different ways. One, it dispatched commands from the plan for execution. Two, it monitored the success or failure of these commands. And three, it monitored a set of resource and state information including items such as rover position, power levels, and onboard memory levels. If unexpected events occurred, then the plan was dynamically revised to accommodate the new information. Note, that in these early tests all unexpected events were undesirable (e.g., resources oversubscribed, traverses taking longer than estimated), and none corresponded to new science opportunities.

The above-mentioned tests will be significantly expanded on this year, including 1) testing all components using real rovers 2) testing the online incorporation of new science goals, and 3) testing with additional real data sets gathered during rover traverses both on earth and Mars.

4 Related Work

The idea of having a scientific discovery system direct future experiments is present in a number of other systems. Work on learning by experimentation, such as IDS [Nordhausen and Langley, 1993] and ADEPT [Rajamoney, 1990], varied certain quantitative and qualitative values in the domain and then measured the effects of these changes. OASIS differs from these systems in that it interacts with the environment to perform experimentation and it is specialized to particular problems and scenarios in planetary science. OASIS is also integrated with a planning system, which constructs the detailed activity sequence needed to perform new science experiments.

Other work has used experimentation to learn from the environment but experiments again have not been scientifically driven. EXPO [Gil, 1993] integrates planning and learning methods to acquire new information by interacting with an external environment. However, while EXPO tries to improve its planning-related domain knowledge, OASIS learns prioritization models of geological terrain features. Another example of learning about the environment is the Minerva museum tour-guide robot [Thrun, et al., 1999]. Minerva learns several pieces of environment knowledge, including maps, sensor models and travel times between museum locations. Minerva's reactive planner uses learned travel times to dynamically alter its tour of the museum based on time limits. Conversely, OASIS learns new prioritized science goals for the rover to achieve, and uses an activity planner/scheduler that reasons about mission goals and resource and state constraints.

Several researchers have addressed methods for extracting features from data with the intention of performing the operations onboard a spacecraft. Gulick *et al.* [2001] presented methods for locating rocks in an image using information about the sun angle, identifying the horizon and rec-

[1] Performance numbers reported in this paper were run on a Linux 1.7 GHz Pentium 4 workstation.

ognizing layers. Gazis and Bishop [Gazis and Bishop, 2002] and Ramsey et al [Ramsey, *et al*., 2002] have both looked at analyzing point spectra for mineral detection. There has also been work on developing a framework for feature extraction and event detection onboard Earth orbiting satellites (EVE) [Tanner, *et al*., 2001]. Our work has specifically focused on identifying and analyzing rocks in gray-scale images thus far and, in contrast to the work mentioned here, takes the next step of using the feature extraction to determine desirable additional actions a rover could autonomously take.

A number of other systems have used planning methods to coordinate robot behavior. [Gat, 1992; Bonasso, *et al*., 1997; Alami, *et al*., 1998]. However, these systems generate plans in a batch approach where plans are generated for a certain time period and if re-planning is required, an entire new plan must be produced. In OASIS, plans are continuously modified in response to changing conditions and goals. The CPS planner, which is also directed towards rover operations, generates contingent plans, which are then executed onboard and can be modified at certain points if failures occur [Bresina, *et al*., 1999]. This planner takes a more limited approach then the OASIS planner, since the only plan modifications that can be performed during execution are those that have previously identified as possible change points. Furthermore, none of these systems has been integrated with a machine learning system that drives future plan goals.

5 Conclusions

This paper presents the OASIS system, which is being developed to support autonomous science operation during long-range rover traverses. OASIS integrates techniques from machine learning with planning and scheduling to dynamically analyze science data, request new science operations, and generate a new plan of action to support those requests. Often, in current rover missions, volumes of data are collected during rover traverses, however much of this data cannot be sent back to earth due to communication restrictions. OASIS enables this data to be analyzed onboard the rover and then used to determine new science measurement goals for objects of high interest. This system is currently being tested using several real rovers and using data gathered during rover field experiments.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[Alami *et al*., 1998] R. Alami, R. Chautila, S. Fleury, M. Ghallab, and F. Ingrand. An Architecture for Autonomy. *International Journal of Robotics Research*. 17(4) April, 1998.

[Bonasso *et al*., 1997] R. Bonasso, R. Firby, E. Gat, D. Kortenkamp, D. Miller, and M. Slack. Experiences with an Architecture for Intelligent, Reactive Agents. *Journal of Experimental and Theoretical Artificial Intelligence Research*, 9(1), 1997.

[Bresina *et al*., 1999] J. Bresina, K. Golden, D. Smith, D., and R. Washington. Increased Flexibility and Robustness of Mars Rovers. *Proceedings of the International Symposium, on AI, Robotics and Automation for Space*. Noordwijk, The Netherlands, June 1999.

[Castano *et al*., 2002] R. Castano, R.C. Anderson, J. Fox, J.M. Dohm, A.F.C. Haldemann, and W. Fink. Automating shape analysis of rocks on Mars. *Proceedings of the Lunar and Planetary Science Conference*, March 2002.

[Castano *et al*., 2003] Rebecca Castano, Robert Anderson, Tara Estlin, Dennis Decoste, Forest Fisher, Daniel Gaines, Dominic Mazzoni, and Michele Judd. Rover Traverse Science for Increased Mission Science Return. *Proceedings of the 2003 IEEE Aerospace Conference*. Big Sky, Montana, March 2003.

[Chien, *et al*., 2000] Steve Chien, Russell Knight, Andre Stechert, Rob Sherwood, and Gregg Rabideau. Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling. *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, April 2000.

[Estlin, *et al*., 2002] Tara Estlin, Forest Fisher, Daniel Gaines, Caroline Chouinard, Steve Schaffer, and Issa Nesnas. Continuous Planning and Execution for a Mars Rover. *Proceedings of the Third International NASA Workshop on Planning and Scheduling for Space*. Houston, TX, Oct 2002.

[Gat, 1992] Erann Gat. ESL: A Language for Supporting Robust Plan Execution in Embedded Autonomous Agents, *Proceedings of the Tenth National Conference on Artificial Intelligence*, San Jose, CA, July 1992.

[Gazis and Bishop, 2002] P. Gazis and J. Bishop. Development of rule-based autonomous spectral analysis techniques for planetary surfaces: preliminary results using lab spectra. *American Geophysical Meeting Fall meeting*, San Francisco, CA, Dec 2002.

[Gil, 1993] Yolanda Gil. Efficient domain-independent experimentation. *Proceedings of the Tenth International Conference on Machine Learning*. 1993

[Gilmore, *et al*., 2000] M. Gilmore, R. Castano, T. Mann, R. C. Anderson, E. Mjolsness, R. Manduchi, and R. S. Saunders. Strategies for autonomous rovers at Mars. in *J. of Geophysical Res*., Vol. 105, No. E12, Dec. 2000 pp. 29223-29237.

[Gor, *et al*., 2001] V. Gor, R. Castano, R. Manduchi, R. Anderson, and E. Mjolsness. Autonomous Rock Detection for Mars Terrain. *Space 2001, American Institute of Aeronautics and Astronautics*, Aug. 2001.

[Gulick, *et al*., 2001] V. Gulick, R. Morris, M. Ruzon, and T. Roush. Autonomous image analyses during the 1999 Marsokhod rover field test. *Journal of Geophysical Research-Planets*, 106 (E4): 7745-7763 Apr. 2001.

[Mishkin, *et al*., 1998] A. Mishkin, J. Morrison, T. Nguyen, H. Stone, B. Cooper, B. Wilcox. Experiences with Operations and Autonomy of the Mars Pathfinder Microrover. *Proceedings of the 1998 IEEE Aerospace Conference*. Aspen, CO, March 1998.

[Nordhausen and Langley, 1993]. B. Nordhausen and P. Langley. An integrated framework for empirical discovery. *Machine Learning*. 12:17-47.

[Ramsey, *et al*., 2002] J. Ramsey, P. Gazis, T. Roush, P. Spirtes, and C. Glymour. Automated remote sensing with near infrared reflectance spectra: carbonate detection. *American Geophysical Meeting Fall meeting*, San Francisco, CA, Dec. 2002.

[Rajamoney, 1990] S. Rajamoney, S. A computational approach to theory revision. In Shrager, J., and Langley, P., eds. *Computational Models of Scientific Discovery and Theory Formation*. San Mateo, CA: Morgan Kaufman. 225-254.

[Stentz, 1994] Anthony Stentz. Optimal and Efficient Path Planning for Partially-Known Environments. *Proceedings of the IEEE International Conference on Robotics and Automation*, San Diego, CA, May 1994.

[Tanner, *et al*., 2001] S. Tanner, K. Keiser, H. Conover, D. Hardin, S. Graves, K. Regner, R. Wohlman, R. Ramachandran, M. Smith. EVE: An on-orbit data mining testbed. *IJCAI-01 Workshop on Knowledge Discovery from Distributed, Heterogeneous, Autonomous, Dynamic Data and Knowledge Sources*. Seattle, Washington, Aug. 2001.

[Thrun, 1999] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. MINERVA: A Second-Generation Museum Tour-Guide Robot. *Proceedings of the IEEE International Conference on Robotics and Automation*, Detroit, MI May 1999.

[Zweben, *et al*., 1994] M. Zweben, B. Daun, E. Davis, and M. Deale. 1994. Scheduling and Rescheduling with Iterative Repair, In *Intelligent Scheduling*, Morgan Kaufmann, San Francisco, CA. 1994. 241-256.